Fermilab

Closeout Presentation

Director's Conceptual Design Review of the LHC CMS Detector Upgrade Project

May 14-16, 2013

This page intentionally left blank

Table of Contents

This page intentionally left blank

Executive Summary

This Independent Design Review, commissioned by the Fermilab Director, is to verify that the US LHC CMS Detector Upgrade Project design is technically adequate to achieve its scientific goals and that it has advanced in maturity to a level appropriate for DOE CD-1.

The LHC will resume operations in 2015 and is expected to exceed the design luminosity by as much as a factor of 1.5 after 1 - 2 years of running. The LHC then plans a long shutdown in 2018, after which will follow at least 3 to 4 years of operation at much higher luminosity, perhaps by a factor of 2 or 3, than it was originally designed to achieve and at which CMS was designed to operate. By 2019 peak luminosities are expected to reach 2-3 x 10^{34} cm-^{2} s^{-1}, corresponding to 50 to 80 interactions per crossing and 25 ns bunch spacing. The CMS detector requires upgrades to the Pixelated Inner Tracking {Pixel} Detector, the Hadron Calorimeter (HCAL) Detector and the Trigger to operate in these conditions.

The LHC CMS Detector Upgrade Project, funded by the DOE and NSF, will upgrade CMS detector components which were provided by the original US CMS detector construction project. The TPC for the upgrade project is currently estimated at ~$43.5M; the NSF is to provide ~1/4 of this funding. The project was granted DOE CD-0 in September 2012, will be reviewed for CD-1 in FY13, and will be complete (CD-4) in the 1st quarter of FY19.

The US CMS Upgrade Project team is a strong, capable team, with experience in the original US CMS detector construction project. At present the US CMS Operations Manager is also serving as the Interim Project Manager. Fermilab should proceed expeditiously to provide a full-time project manager to lead this project. This team has a good understanding of the existing CMS detector and the improvements necessary to operate efficiently while maintaining physics capability in conditions that will exist after the long shutdown in 2018. Science goals and physics requirements have been documented and translated into detector requirements and specifications; the designs presented will likely meet these requirements and specifications. Designs can be constructed, inspected, tested, installed, operated and maintained; and prototypes exist for many of the designs. Documentation in support of the conceptual design is extensive and will support advancing to the next stage of design following CD-1. The designs are in all cases judged to be well advanced and in most cases are beyond the conceptual design level. They provide a good basis for establishing the cost and schedule range and are therefore a good foundation for proceeding to DOE CD-1.

1.0 Introduction

A Director's Conceptual Design Review of the LHC CMS Detector Upgrade Project was held on May 14-16, 2013 at the Fermi National Accelerator Laboratory. The object of this review was to assess the status and adequacy of the overall CMS Detector Upgrade conceptual design effort to meet the requirements for a DOE Critical Decision 1 (CD-1) "Approve Alternative Selection & Cost Range". The charge included a list of topics and specific questions to be addressed as part of the review. The assessment of the Review Committee is documented in the body of this closeout presentation.

Each section in this closeout presentation is generally organized by Findings, Comments and Recommendations. Findings are statements of fact that summarize noteworthy information presented during the review. The Comments are judgment statements about the facts presented during the review and are based on reviewers' experience and expertise. The comments are to be evaluated by the project team and actions taken as deemed appropriate. Recommendations are statements of actions that should be addressed by the project team. The remainder of this presentation has the answers to the review charge questions.

The LHC CMS Detector Upgrade Project is to develop a response to the review recommendations and present it to the Laboratory Management and regularly report on the progress during the Project's Project Management Group Meetings (PMGs) and at the Performance Oversight Group (POG). The recommendations will be tracked in the iTrack system where progress to closure will be tracked.

2.0 Hadron Calorimeter - HCAL

Contributors: Dmitri Denisov, Adam Para

Findings

- Current photodetectors, HPD's, in HB/HE CMS calorimeters are not adequate to reach required CMS physics performance in the future high luminosity running. HPD detectors are no longer produced and their replacement by a novel type of photodetectors, SiPM's was presented as the only practical option for the CMS HCAL upgrade.

- Superior performance and segmentation of SiPMs allows for a split of the hadron calorimeter towers into several depth segments. The proposal is to have three depth segments in HB and five in HE. A replacement of the photodetectors including depth segmentation can be achieved without changes to the calorimeter absorbers, scintillators, and light guides and are limited to the replacement of a "Readout Boxes" containing the Optical Decoder unit, SiPM's and their control card, and front-end electronics cards.

- Long process of R&D and collaboration with the vendors has demonstrated that the SiPM's provide adequate performance, including long term stability and radiation resistance, for the high luminosity running of the CMS. Two vendors have been identified to be able to deliver photodetectors with required parameters and another vendor is close to meet specifications as well.

- Large gain provided by the SiPM's and longitudinal segmentation of the calorimeter necessitate the replacement of the front-end electronics. A new generation of the QIE ASIC (versions 10 and 11) has been developed for this purpose. The number of readout channels is increased by a factor of about three to allow independent readout of the longitudinal segments. The QIE 10/11 provide a new functionality: signal leading edge timing information with the 0.5 nsec bins. The design and testing of the new electronics is advanced, approaching a pre-production phase.

- SiPMs will be calibrated using LED calibration system. A HB/HE Calibration Module conceptual design exists and is based on re-using the infrastructure of the QIE cards.

- New front-end electronics modules, including the upgrade of the fiber links required by increased data rates, new generation Clock Control Module and increased functionality, have been designed and prototyped.

- Increased channel count and functionality results in the increase of the power dissipated by the Readout Module by a factor of three. It is expected that the existing water cooling system will be sufficient for the operation, albeit at the temperature increase between input and output of the module by 3 degrees C.

- Detector monitoring and safety systems will be upgraded to provide adequate control and monitoring of the new Readout Modules and to preserve safety of the remote detector elements.

- New back-end electronics microTCA standard modules allowing for the increased data transfer rates and including new clock control modules have been designed and prototyped.

- Cherenkov light produced in the windows of PMTs of the HF were found to be a source of anomalous signals in the forward calorimeter. This is being remedied by replacing a single PMT tube with multianode PMT configured to provide dual output. Adapter boards have been designed and prototyped to adapt this dual output to the present configuration of the CMS detector.

- Reduction or elimination of the spurious signals in HF will be realized with the new upgraded front-end and back-end readout electronics, including timing information provided by the QIE10 card. Front-end and back-end readout cards for all CMS calorimeters are functionally similar and use many of the same components. They differ in packaging due to different spatial constraints.

- Optical splitters for back-end electronics have been designed and prototyped to provide parallel data path for the trigger for post LS1 running. These splitters are mounted on the HB/HE electronics racks and necessary for the development and testing of the upgraded trigger. Use of the splitters will lead to a significant, but acceptable, reduction of the light signals.

Comments

- CMS HCAL upgrade project presented for the review is well advanced with main technical solutions well described and documented. In many cases working prototypes exist, work is well advanced on technical developments and simulation. The project is advanced to beyond "conceptual design" stage.

- Technical solutions proposed are scientifically and technically sound and are expected to satisfy long term successful operation of the CMS HCAL.

- Main parts of the calorimeter, including absorbers, scintillators and light guides are not planned to be replaced while photodetectors and electronics (both front-end and back-end) undergoing substantial upgrades. Ability of the calorimeter elements which are not planned to be upgraded to perform according to the specifications looks adequate for the goals of the upgrade while radiation and other potential sources of aging have to be carefully monitored.

- Proposed upgrade is expected to provide substantial improvements to the calorimetry by providing longitudinal towers segmentation and timing of the signals arrival. Software developments to use this substantially increased amount of information to improve calorimeter performance, including particle flow algorithms, have to progress expeditiously.

- SiPMs is the heart of the HCAL upgrade project, except very forward part of the calorimeter where PMTs are used. These are revolutionary new devices for light detection with many unique properties well matched to the HCAL upgrade goals. At the same time rapid development of these devices by manufacturers might introduce unexpected features which could be discovered later in the project or during CMS high luminosity data collection. Selecting final SiPM type and vendor should be done at appropriate time, so all properties of these new photodetectors could be carefully investigated.

- Cost estimates are in most cases reasonable and based on extensive past experience of the CMS HCAL team and direct quotes from vendors. Number of proposed spares is adequate.

- Preliminary schedule looks achievable while with many groups/vendors involved coordination of the project will require substantial efforts and flexibility.

- The project is based on the experience of the CMS HCAL construction, installation, commissioning and operation with many key experts involved in the HCAL for a long time. This adds assurance that the project will be well integrated into CMS and reduces technical risks.

- Due to the size of the project there are quite a few risks involved. The project team provided preliminary estimates of the most vulnerable areas, while careful analysis of risks and especially ways to mitigate them has to be among project priorities.

- With many deliverables distributed over extended period of time the project has to develop optimal schedule to assure technical availability of parts at the time they are actually needed while fitting into available manpower and funding profiles.

- There are CMS HCAL improvements which are already under implementation/construction or have to be within about a year in order to fit into overall LHC/CMS schedule. Clear identification of the HCAL improvements which "belong to the project" will help with crisp definition of the project scope.

- During the review HCAL upgrade participants demonstrated deep understanding of the technical details and were able to answer on multiple questions well. The team has deep understanding of the project goals, ways to achieve them as well as possess strong experience based on the previous involvement in CMS HCAL.

- Project management team is professional and deeply familiar with all elements of the project including technical details, manpower, costs and schedule. US CMS HCAL upgrade team has to be congratulated with developing technically sound project, providing in depth documentation and excellent talks for the review.

- CMS HCAL upgrade project is well developed for CD-1 review in August 2013.

Recommendations

1. Develop a list of all major upgraded hadron calorimeter performance parameters driven by CMS physics goals.

2. Document the procedure for the absolute calibration of the calorimeter energy scale, including amount of data and specialized triggers required to collect samples adequate for precision calibration.

3. Develop a plan for how to finalize the decision on SiPMs selection and procurement and on in depth tests of the selected SiPMs performance, including long term stability of operation.

3.0 Silicon Pixel Detector - FPIX

Primary Writer: Marina Artuso

Contributors: Dave Christian

3.1 General Remarks

Findings

- The upgraded pixel detector will tolerate much higher rate without loss of data than the current detector.

- The upgraded pixel detector includes 4-barrel layers and 3 forward disk layers at each end of the barrel (6 forward disks in total).

- The upgraded detector will have lower mass in the tracking volume than the current detector.

- The existing installed cable plant will be used for power distribution, control, and readout of the upgraded detector. The upgraded detector will require more power than the existing detector. Use of the existing cable plant will be enabled by the use of DC-DC converters in the detector volume and 10V distribution from power supplies to the detector. Existing dark fibers will be used for readout.

- The project goal is to have the detector ready for installation in the experimental area during the time between LS1 and LS2. Installation before LS2 will require an extended technical stop.

Comments

- Increased rate capability is essential to the efficient operation of sections of the pixel detector closest to the circulating beams at high luminosity.

- The increase in the number of pixel layers is very well motivated by Monte Carlo studies of tracking in the presence of the projected number of pileup events.

- Most aspects of the project are far beyond the conceptual design level and are close to be ready for production.

- The detailed design of half disks and the development of half disk fabrication procedures are less advanced than other aspects of the project.

Recommendations

None.

3.2 Detector Module

Findings

- The detector modules will use the same technology as the current detector and will be produced by the same vendor.
- The new detector will be constructed entirely with a single module configuration.
- Bump bonding will be done by one or two commercial vendors, well known to the CMS group.
- The fabrication of detector modules from bump bonded sensor-readout chip hybrids and high density interconnect circuits will be performed at two assembly sites, using identical robotic assembly procedures.

Comments

- A design using one module type simplifies detector design and fabrication and is recognized as a "best practice." It also significantly reduces the risk associated with sensor and bump bonding yield compared to a design requiring many sensor types.

Recommendations

None.

3.3 Electronics

Findings

- All ASICs will be fabricated in a common 0.25μm CMOS process.
- The pixel readout chip (ROC) will be supplied by PSI.
- The most significant ancillary ASIC will be a new "Token Bit Manager" (TBM).
- The detector will use a new power distribution scheme based on CERN supplied DC-DC converters.
- One new type of HDI is required; a prototype has been fabricated.
- New "port-cards" will be located in the pixel service cylinder and will translate high-speed output data from electrical to optical signals.
- A new Al flex cable will carry power, HV, and signals between pixel modules and port-cards. Prototypes have been produced by two vendors.
- Plans for the replacement of "Front End Digitizers" (FEDs) have not been finalized (two solutions are under study).

Comments

- The current version of the ROC may have sufficient rate capability for FPIX, and is almost certain to be useable for the pilot system.
- Early engineering readiness reviews of the ROC and TBM are planned. These will be important milestones for the development of a realistic module production schedule.
- A fully functional DC-DC converter is not yet available. The current version is close enough to fully functional so that it can be used in the pilot system if necessary.
- Module tests will be important for the validation of the HDI design (analog and digital performance).
- System tests could play an important role in the characterization of the Al flex cable.
- We strongly support the timely selection of a single solution for the FED board replacement in order to avoid duplication of effort.

Recommendations

None.

3.4 Mechanical structure and cooling

Findings

- The new half disk mechanical and cooling design is significantly different from that of the installed detector.
- The new half disk is divided into inner and outer sectors with different orientation of the sensors to facilitate optimal charge sharing among pixels.
- The new detector will use CO_2 cooling and operate at lower temperature than the existing detector.
- The service cylinders for the new detector will be very similar to the existing service cylinders. Installation procedures and fixtures will require only minor modifications.
- Current plans call for the construction of twelve half disks and four support cylinders.

Comments

- Mechanical and cooling design is at the level of an advanced conceptual design.
- Consider the construction of spare half disks (without pixel modules).

Recommendations

None.

3.5 System assembly and testing

Findings

- Half disks will be built as mechanical structures and tested with heaters before sensor modules are placed.

- The module mounting procedure is still being developed.

- Half disk testing after assembly will include a full test with all modules powered and reading out simultaneously.

Comments

- We strongly suggest planning system tests of increasing scale rather than relying exclusively on the pilot system and on tests of completed half disks. In particular, test beam runs of "slices" of the full system are very likely to be instrumental to an optimization of the system design.

Recommendations

None.

3.6 Pilot system

Findings

- Two partial half disks, each with two blades carrying two pixel modules (for a total of 8 pixel modules) will be installed in the current FPIX during LS1. These half disks will be configured as closely as possible to the final FPIX upgrade design and will be read out into the CMS DAQ during next LHC running period. One of the two partial half disks will use DC-DC converters.

Comments

- This system will provide very useful operating experience and be a great asset for software development.

Recommendations

None.

4.0 Level 1 Trigger

Contributors: Bill Ashmanskas, Eric James, and Peter Wilson

Findings

- In the absence of trigger upgrades, the increased pile-up anticipated for LHC operation after LS1 will cause degraded Level 1 trigger performance, such as reduced background rejection at a given efficiency for isolated-lepton triggers.

- In the period after LS1, with the LHC operating for the first time at full beam energy, it will be especially important for the CMS trigger to preserve its sensitivity to a wide range of physics processes.

- To increase the trigger's background rejection while preserving its signal efficiency, the proposed upgrade makes two key modifications: more detailed information per bunch crossing is transmitted from the detector to the trigger; and vastly increased computational power is provided in the trigger system for processing that information. The upgraded processing power takes the form of high-connectivity state-of-the-art FPGAs, which offer considerable flexibility for CMS to adapt its trigger algorithms as run conditions evolve and as LHC physics experience is gained.

- The proposed trigger upgrade will use state-of-the-art telecommunications technology to support the increased bandwidth requirements imposed by the higher granularity of calorimeter data and by the larger number of Local Charged Track muon-chamber segments transmitted from the collision hall to the calorimeter and muon trigger systems, respectively.

- The project team includes the same people who designed and commissioned the original CMS trigger hardware. Team members' collective experience spans several other major HEP trigger systems. The team shows a seasoned understanding of the process of making a new trigger system work and presented a detailed plan for a phased implementation of the proposed upgrade.

- The hardware design includes many handles to facilitate rapid testing and commissioning, such as test-pattern injection, spy-buffer readout, test stands, and facilities for vertical-slice tests.

- The design emphasizes common hardware choices among subsystems, e.g. uTCA infrastructure, the Virtex-7 FPGA, and multi-gigabit optical links, which allow the upgrade team to use experts' time efficiently, to avoid redundant effort, and to develop new hardware modules quickly.

- Working prototypes exist for nearly all hardware components of the new design. The Virtex-7 FPGAs used in the new trigger-processing modules are available for purchase, and the proponents have vendor quotes for the Virtex-7 in hand.

- The connectivity and processing power of the upgraded calorimeter trigger hardware can support a range of event-processing methodologies. Two such methodologies were presented: a traditional, pipelined approach in which distinct modules process distinct regions of the detector; and a novel “time-multiplexed trigger” (TMT) approach in which each of N distinct modules processes the complete calorimeter’s data for every Nth event. Though the relative merits of the two approaches are still under investigation, the CMS collaboration presently favors the TMT approach as its preferred option.

- Throughout the design process, the proponents have already carried out alternative analysis and value engineering, for example to select FPGA families and crate/chassis infrastructure. The CDR draft available at the time of this review does not document the proponents’ alternative analysis and value engineering explicitly.

- The design of upgraded trigger systems is constrained by limitations to CMS that will persist until at least LS3 but are not documented in the CDR. For example, the Level 1 accept rate is limited to 100kHz by the DAQ and High Level Trigger and the total L1 processing time is limited by front-end buffer depth.

- Both the calorimeter and muon trigger upgrades are based on very powerful FPGAs as processing engines. In addition to the very mature hardware development projects for this hardware there will be substantial firmware and software efforts to develop trigger algorithms, configuration, debugging, and monitoring tools. The planned effort on the project for the firmware alone is approximately 20 FTE-years.

- The proponents presented a plan to parasitically commission the upgraded trigger system in parallel with the existing trigger during physics data taking starting in 2015. The plan includes implementing increasing functionality in stages over a period of several years. As increased functionality becomes available it could be included in the trigger decision. Some components of this plan would involve funding from the Operations program in addition to the Upgrade project.

- To provide input data to the upgraded trigger hardware in parallel with the existing system, either signals must be split passively (HCAL) or new mezzanine cards that drive the data to both systems must be built. These cards plus the interconnect fibers must be installed during the current long shutdown (LS1) to meet the staged commissioning plan. The Muon Port Card (MPC) used for the CSC system has been designed, prototyped, and tested. It will be ready for the production by Q1 of FY14, which is necessary for installation during LS1.

Comments

- The committee agrees with the proponents that the design architecture (high input bandwidth and flexible processing) is ideally matched to the CMS physics requirements.

- The choice of telecommunications industry standards for transmitting information between boards is a good one.

- Because of the significant previous experience of the project team members, we have high confidence in their ability to implement the proposed upgrades.

- The current state of the development (having working prototypes in hand for essentially all of the components) gives us high confidence in the success of the project. As a result, we consider the technical risks associated with the project to be low.

- The case for the chosen architecture of the upgraded calorimeter trigger is strengthened by, but in no essential way depends on, the choice by the CMS collaboration of a novel time-multiplexed trigger (TMT) as its preferred option. Presenting the TMT option as the project baseline may be unnecessary (and an additional complication) for the DOE review process.

- The proponents should be commended for the excellent work examining possible algorithms with the potential to provide the required rate reductions and efficiency improvements. For example, the possibility of implementing multi-variate techniques for muon track finding looks quite promising. Since these are only example algorithms, it is critical that a high degree of flexibility be maintained in the hardware and the importance of that flexibility be communicated to future review teams clearly.

- Prior to CD-1 review the CDR should be strengthened to reflect the alternative analysis and value engineering that has already been carried out in the design process.

- More information about the specific constraints of the trigger system (Level 1 latency and accept rate) and why these are fixed within the context of this upgrade should be provided in the CDR.

- The firmware and software plans presented to the committee seemed well thought-out. Resource needs should be to be further fleshed-out prior to the upcoming cost and schedule reviews for CD-1.

- The incremental approach to the trigger upgrade makes good technical and strategic sense, and necessarily involves a mix of operations and project funding. For the CD-1 cost and schedule review, however, the boundary between operations and project should be described more clearly.

- Installing the new Muon Port Cards prior to the end of LS1 is important, but it is not clear how this can be achieved within constraints of upgrade project funding.

Recommendations

4. Develop a plan prior to the CD-1 review for procurement of components that need to be installed prior to the end of LS1.

5. Incorporate the missing information identified within comment section into the CDR prior to the CD-1 review.

5.0 Charge Questions

5.1 Are the science goals and physics requirements clearly stated and documented? Have the science goals and physics requirements been adequately translated into technical performance requirements and specifications?

HCAL – Yes. The proposed CMS HCAL upgrade is intended to maintain the present excellent performance of the CMS calorimeter for the high luminosity running. Physics requirements for studies of Higgs boson properties, searches for new physics, and precision measurements of the Standard Model parameters are translated into technical requirements and specifications for the elements of the project well. As this project is an upgrade of the existing detector technical specifications in many cases driven by the requirement of compatibility with geometrical and other constraints of the existing detector.

FPIX – Yes. The physics case for the FPIX upgrade is compelling: the higher efficiency and lower fake rate facilitated by the improved system design will have a strong impact on the offline reconstruction efficiency and the optimization of the high level trigger. The detailed requirements for individual project subsystems to achieve the high level goals could be made more specific.

Trigger - Yes. The basic goal of maintaining current trigger performance at higher instantaneous luminosity and pile-up is clearly stated. High efficiency at low thresholds will be particularly important for both studies of the Higgs boson and searches for potential new physics signals.

5.2 Is the design technically adequate? Is the design likely to meet the technical requirements needed to carry out the scientific goals?

HCAL – Yes. The proposed design is technically adequate for the CMS HCAL upgrade and likely to meet the objectives.

FPIX – Yes. Several components of this project have reached a level of maturity and planning far exceeding the scope of our review. Some items related to the mechanical design are in a more conceptual phase that is rapidly evolving toward a design that meets all the requirements to carry out the scientific goals.

Trigger - Yes. Convincing evidence was shown that the current trigger menu can be maintained at similar or better levels of performance with the proposed upgrades.

5.3 Can the design be constructed, inspected, tested, installed, operated and maintained in a satisfactory way?

HCAL - Yes, the design has been worked out in a considerable depth, beyond the stage of a conceptual design. The readiness of the design in many cases is close to a pre-production stage.

FPIX – Yes. This project has solid foundations. Several elements are constructed extrapolating from the excellent experience in the construction and installation of the FPIX detector currently operating in CMS. We are confident that the proponents will be successful in translating this design into an excellent system that can be installed, operated, and maintained in a way that will fully meet all the requirements.

Trigger - Yes. Considerable thought has gone into designing upgrades that can be plugged in with minimal disruption to the running experiment. The team is very experienced, having been involved in the production of much of the original hardware, and a solid plan exists for testing and commissioning the hardware.

5.4 Is there adequate supporting documentation to support the conceptual design and the transition to developing the preliminary design?

HCAL – Yes. There is an extensive documentation available, including technical details. The breadth and depth of the available documents exceeds the capacity of a small group of reviewers and limited time available for the review.

FPIX – Yes. Most elements of the design are well beyond the conceptual level.

Trigger - Yes. The team demonstrated a design whose maturity is well beyond the conceptual stage and in many cases past the preliminary stage. This maturity was demonstrated through a combination of material in breakout presentations and the formal CDR.

5.5 Are the risks (on technical, cost, and schedule basis) of the selected base design approach and alternatives understood and are appropriate steps being taken to manage and mitigate these risks? Have areas been identified where value engineering should be done? If value engineering has been performed is it documented?

HCAL - Various risks have been identified and options to mitigate them presented to the reviewers. The primary risks are related to the availability of the experts and/or technical processes required for the fabrication by vendors. Execution of the project without delays is the primary mitigation technique.

FPIX - The identified risks are mostly related to the fact that some key electronics components are custom made and developed by others (CERN, ETH). Some consideration should be given to risks associated with the novel aspects of the mechanical and cooling design. One example of value engineering is the ongoing effort to identify an optimum solution for the replacement of the FEDs.

Trigger - Yes. The risks associated with the cost and production of materials is low due the advanced stage of the hardware designs. Firmware and software contributions to the project are well defined but carry a somewhat higher risk because this effort is at a more preliminary stage. Alternatives analysis and value engineering have been performed during the course of the design process but should be better documented.

5.6 Are the project organization and lines of responsibility clearly defined and sufficient to ensure the successful engineering and design of the project?

Yes, the project is organized in a manner consistent with the WBS structure and therefore along lines of deliverables. Lines of responsibility and authority are clearly defined in the PEP.

Are the design interfaces between the US CMS Upgrade Project and the International CMS Upgrade at CERN understood and well defined to ensure a coordinated effort and an integrated design?

Design interfaces are well understood; protocols and specifications are documented. The CMS Technical coordinator, before construction, reviews the soundness and completeness of designs, including the coherence of all interfaces.

Is there a reasonable plan in place for implementing configuration management to ensure changes to the technical requirements/specifications are controlled and communicated to all affected groups?

There is a configuration plan under development to employ change control boards and a capable document management system. This plan will be in place by CD-1.